

What is claimed is:

Sub
A11
1. A method for recovering from failures affecting a
resource manager within a group of resource managers,
5 wherein the resource managers within the group have access
to a shared resource via which remote resource managers
communicate with the resource managers within the group,
the shared resource including data storage structures to
which resource managers within said group connect to send
10 and receive communications, the method comprising:

storing, within a first data storage structure of the
shared resource, unit of work descriptors for operations
performed in relation to said shared resource by the
15 resource managers in said group;

20 sending a notification of a connection failure between
a second data storage structure of the shared resource and
a first resource manager within said group, the
notification being sent to the remaining resource managers
within the group which are connected to the second data
storage structure;

25 one or more of said remaining resource managers
accessing said first data storage structure and analysing
the unit of work descriptors to identify the units of work

relating to the second data storage structure that were being performed by the first resource manager when the connection failure occurred; and

5 said one or more remaining resource managers recovering the identified units of work.

10 2. A method according to claim 1 wherein, if there are no remaining resource managers connected to the second data storage structure after said connection failure, said notification is sent to a remaining resource manager when that resource manager connects to the second data storage structure.

15 3. A method according to claim 1 wherein, if there are no remaining resource managers connected to the second data storage structure after said connection failure, the failed resource manager determines when it is restarted whether any other resource manager has performed recovery for its
20 units of work relating to the second data storage structure and, upon determining that no resource manager has performed said recovery, the restarted resource manager recovers said units of work.

25 4. A method according to claim 1, wherein all remaining resource managers within the group which are connected to

the second data storage structure respond to said notification by attempting to access said first data storage structure to identify units of work to recover, and the method includes the further steps of:

5

responsive to a first remaining resource manager identifying a unit of work to recover, said first remaining resource manager attempting to set a flag for said unit of work;

10

responsive to successfully setting said flag, assigning recovery responsibility for said unit of work to said first remaining resource manager; and

15

refusing to assign recovery responsibility for said unit of work to said first remaining resource manager if said flag has been set by another remaining resource manager.

20

5. A method according to claim 4, including the further step of:

25

responsive to said flag having been set by another remaining resource manager, said first remaining resource manager attempting to identify a further unit of work to

recover and attempting to set a flag for said identified further unit of work.

6. A method according to claim 4, including the following steps in response to a connection failure between the second data storage structure of the shared resource and said first remaining resource manager during recovery of said unit of work:

sending a notification of said connection failure to the remaining resource managers within the group which are connected to the second data storage structure;

one or more of said remaining resource managers accessing said first data storage structure and analysing the unit of work descriptors to identify the units of work relating to the second data storage structure that were being performed by the first remaining resource manager when the connection failure occurred; and

said one or more remaining resource managers recovering the identified units of work

7. A method according to claim 1, wherein the unit of work descriptors include:

a unit of work identifier;

an identification of messages put or retrieved within the unit of work;

5

a status for the unit of work; and

a sequence number.

10 8. A method according to claim 1, wherein the shared
resource is a coupling facility list structure, the second
data storage structure is a coupling facility list
structure in which a coupling facility list header
represents a shared access message queue, and the first
15 data storage structure is an administration list structure
of the coupling facility for storing unit of work
descriptors.

20 9. A method according to claim 8, including storing
within the coupling facility, for each resource manager
within the group, a list header information map
representing the set of shared access message queues within
the second data storage structure for which the resource
manager has performed some work.

25

10. A method according to claim 9, including reading said list header information map during recovery to identify the set of shared access message queues within the second data storage structure for which the failed resource manager has performed some work.

11. A method according to claim 1, including storing within the shared resource a structure interest map identifying the set of data storage structures to which respective resource managers within said group are connected.

12. A method according to claim 11, wherein the step of recovering the identified units of work is a first recovery phase and wherein the method includes a second recovery phase comprising the steps of:

reading the structure interest map for the failed resource manager to identify the set of data storage structures to which the failed resource manager was connected at the time of said connection failure;

identifying any operations performed by the failed resource manager on said set of data storage structures which were not recovered in the first recovery phase; and

one or more of said remaining resource managers then backing out said unrecovered operations.

13. A method according to claim 12, wherein the method includes setting a key for operations performed in relation to the shared resource, the key identifying the resource manager which performed the operation, and wherein the identification of operations performed by the failed resource manager comprises checking said keys for unrecovered operations performed in relation to any of said set of data storage structures.

14. A method according to claim 1, wherein a single unit of work represented by a unit of work descriptor may include operations performed in relation to a plurality of data storage structures, and wherein the partial units of work corresponding to said operations are recovered by different ones of said remaining resource managers within the group.

15. A method for recovering from failures affecting a resource manager within a group of resource managers, wherein the resource managers within the group have access to a shared resource, the shared resource including data storage structures to which resource managers within said

group connect to perform operations in relation to data held in said shared resource, the method comprising:

5 storing, within a first data storage structure of the shared resource, unit of work descriptors for operations performed by the resource managers in said group in relation to data held in said shared resource;

10 sending a notification of a connection failure between a second data storage structure of the shared resource and a first resource manager within said group, the notification being sent to the remaining resource managers within the group which are connected to the second data storage structure;

15 one or more of said remaining resource managers accessing said first data storage structure and analysing the unit of work descriptors to identify the units of work relating to the second data storage structure that were
20 being performed by the first resource manager when the connection failure occurred; and

25 said one or more remaining resource managers recovering the identified units of work.

16. A method according to claim 15, wherein the data storage structures of said shared resource include data storage structures which contain shared message queues and said operations performed in relation to said shared resource include putting messages onto a shared message queue and retrieving messages from a shared message queue, for communication between a remote resource manager and resource managers within said group.

17. A method according to claim 16, wherein the unit of work descriptors include:

a unit of work identifier;

an identification of messages put or retrieved within the unit of work;

a status for the unit of work; and

a sequence number.

18. A method according to claim 16, wherein the operations of putting messages onto a shared queue and retrieving messages from a shared queue are performed under transactional scope such that a message which is put is only available to resource managers other than the resource

manager putting the message after commitment of the put
operation and a message which is retrieved is only
available to the retrieving resource manager after
commitment of the retrieval operation, and wherein said
5 stored unit of work descriptors identify each of the
following:

10 units of work that were uncommitted but for which a
decision to commit had been made when the failure
occurred;

15 units of work that were uncommitted but for which a
decision to abort had been made when the failure
occurred; and

20 units of work for which no commit or abort decision
had been made when the failure occurred;
and wherein recovering the identified units of work
comprises:

committing message put and retrieval operations for
which a decision to commit had been made;

25 backing out message put and retrieval operations for
which a decision to back out had been made; and

backing out message put and message retrieval operations for which no commit or abort decision had been made.

5 19. A distributed data processing system including:

a plurality of resource managers;

10 a shared access resource including data storage structures to which the resource managers connect to send and receive communications to and from remote resource managers, the shared access resource including:

15 means for storing, within a first data storage structure of the shared resource, unit of work descriptors for operations performed in relation to said shared resource by the resource managers in said plurality; and

20 means for sending a notification of a connection failure between a second data storage structure of the shared resource and a first resource manager within said plurality, the notification being sent to the remaining resource managers within the plurality which

25 are connected to the second data storage structure;

wherein said remaining resource managers include:

means for accessing said first data storage structure and analysing the unit of work descriptors to identify the units of work relating to the second data storage structure that were being performed by the first resource manager when the connection failure occurred; and

means for recovering the identified units of work.

20. A computer program product comprising program code recorded on a machine-readable recording medium, the program code comprising the following set of components:

a plurality of resource managers;

a shared access resource manager including program code for managing storage and retrieval of data within data storage structures to which the resource managers connect to send and receive communications to and from remote resource managers, the shared access resource manager including:

means for storing, within a first data storage structure of the shared resource, unit of work

descriptors for operations performed in relation to said shared resource by the resource managers in said plurality; and

5 means for sending a notification of a connection failure between a second data storage structure of the shared resource and a first resource manager within said plurality, the notification being sent to the remaining resource managers within the plurality which are
10 connected to the second data storage structure;
wherein said remaining resource managers include:

15 means for accessing said first data storage structure and analysing the unit of work descriptors to identify the units of work relating to the second data storage structure that were being performed by the first resource manager when the connection failure occurred; and

20 means for recovering the identified units of work.